EXHIBIT 32

1994 MAGIC Technical Symposium

August 17-18, 1994
Nichols Hall
Telecommunications and Information Sciences Laboratory
Department of Electrical Engineering and Computer Science
University of Kansas
Lawrence, Kansas

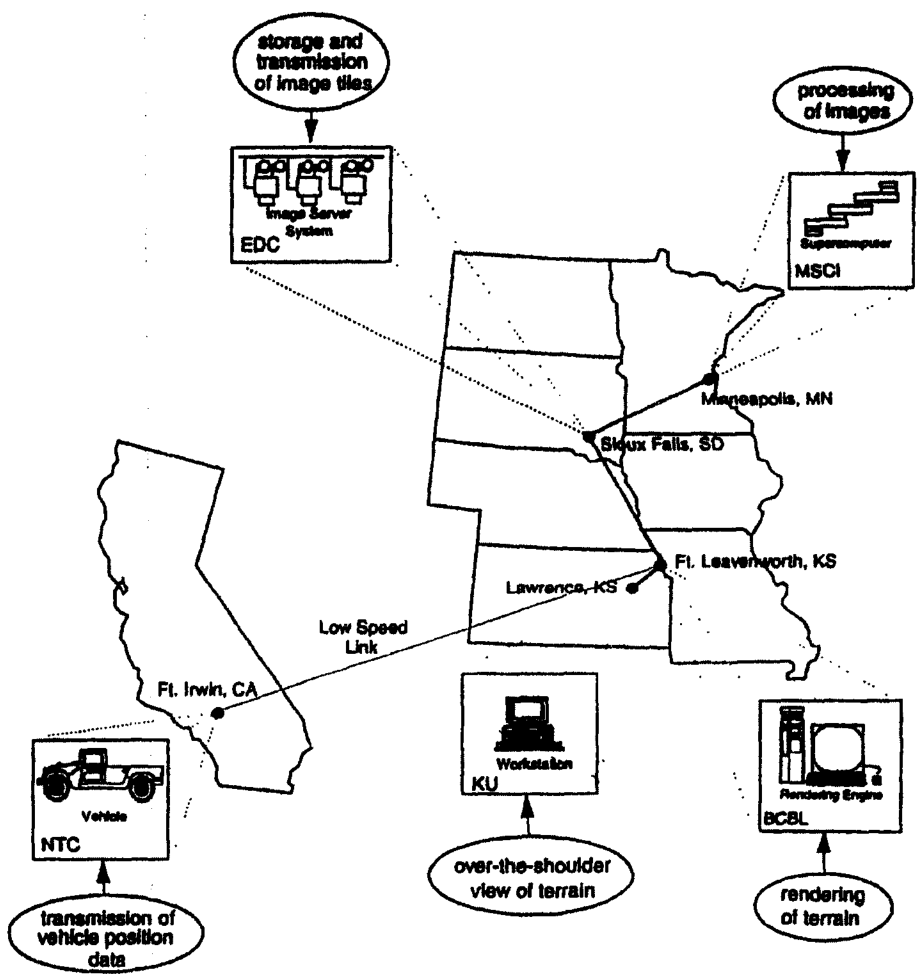

S. LAU
EXHIBIT NO. 201
6-22-06

GOOG0026931

1994 MAGIC Technical Symposium

August 17-18, 1994
Nichols Hall
Telecommunications and Information Sciences Laboratory
Department of Electrical Engineering and Computer Science
University of Kansas
Lawrence, Kansas

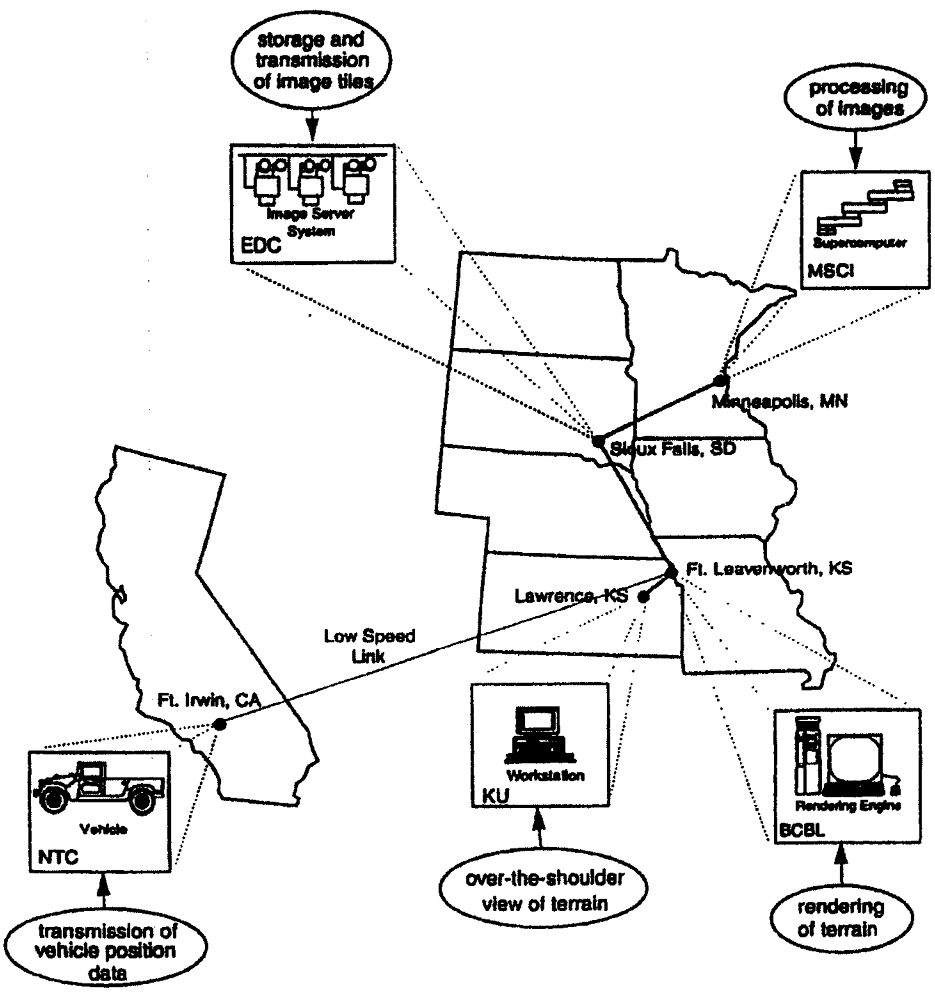

Table of Contents

13. Data Preparation and TerraForm
Jay Feuquay, EDC

14. HIPPI/ATM Gateway: Architecture and Design
Tim Salo, MSC

15. Architecture and Implementations of a 622 Mbps ATM/SONET LAN/WAN Gateway
Victor Frost, KU

MAGIC Technical Papers

1. TCP/ATM Experiences in the MAGIC Testbed
B. Ewy, J. Evans, V. Frost, G. Minden, KU

2. An ATM WAN/LAN Gateway Architecture
J. Evans, V. Frost, G. Minden, D. Petr, KU

3. A Parallel SONET Scrambler/Descrambler Architecture
J. Evans, G. Minden, S. Seetharam, KU

4. ATM Cell Delineation Architectures
G. Gupta, J. Evans, G. Minden, KU

5. A Pipelined Architecture to Map ATM Cells to 622 Mb/s SONET OC-12 Payloads
J. Evans, G. Minden, S. Seetharam, KU

6. A Single Chip Implementation of Receive Path Termination for SONET OC-12c and Quadruple SONET OC-3c
J. Evans, G. Minden, S. Seetharam, KU

7. The Need For Efficient Flow Control Schemes In High Speed Wide Area Networks
J. Evans, G. Minden, S. Seetharam, KU

8. A High Speed Implementation of Adaptive Shaping for Dynamic Bandwidth Allocation
C. Braun, et. al., KU

9. The Image Server System: A High-Speed Parallel Distributed Data Server
   Brian Tierney, William Johnston, Hanan Herzog, Gary Hoo, Guojun Jin, Jason Lee, LBL

10. TerraVision: A Terrain Visualization System
    Yvan Leclerc, Stephen Lau, SRI

11. Protocol Overhead in IP/ATM Networks
    John Cavanaugh, MSC

12. Internetworking with ATM WANs
    John Cavanaugh, Tim Salo, MSC

TerraVision:
A High Speed Terrain Visualization System

Yvan G. Leclerc, Stephen Q. Lau, and Bryan Gorman

What is TerraVision?

‡ TerraVision is a system for visualizing terrain data comprising high-resolution aerial images and elevation data. It allows a user to do the following in real time:

- Pan and zoom over the terrain imagery in 2D
- "Fly" or "drive" over the terrain in 3D.

‡ The aerial images and elevation data are precisely aligned with each other and with a world coordinate system. This enables

- The superposition of buildings and moving vehicles with known real-world coordinates
- The registration of the user's viewpoint to a map
- The specification of new viewpoints by pointing on a map.

How Is the Graphic Image Rendered?

- ‡ Aerial images are preprocessed (orthorectified) so as to be in precise alignment with elevation data.
- ‡ A mosaic of the overlapping orthoimages is created.
- ‡ The orthoimage mosaic is then "draped" (texture mapped) over the terrain during fly-throughs.

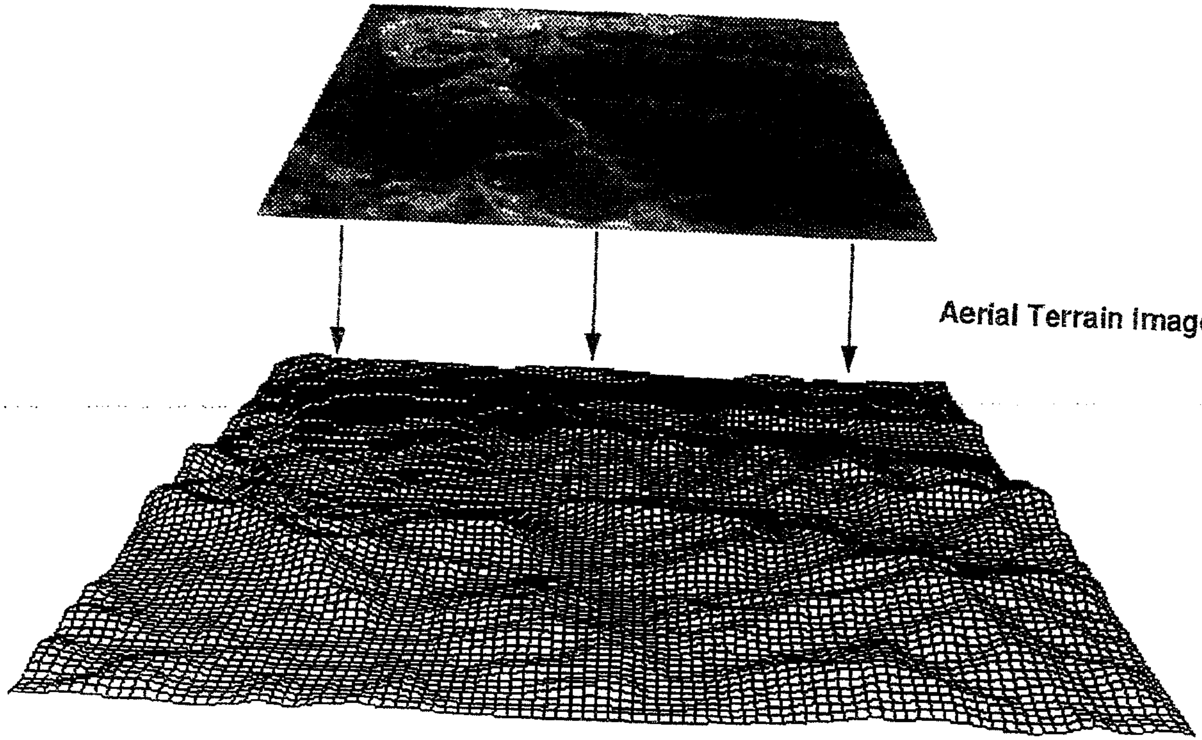

Aerial Terrain Image is Mapped onto Elevation Model

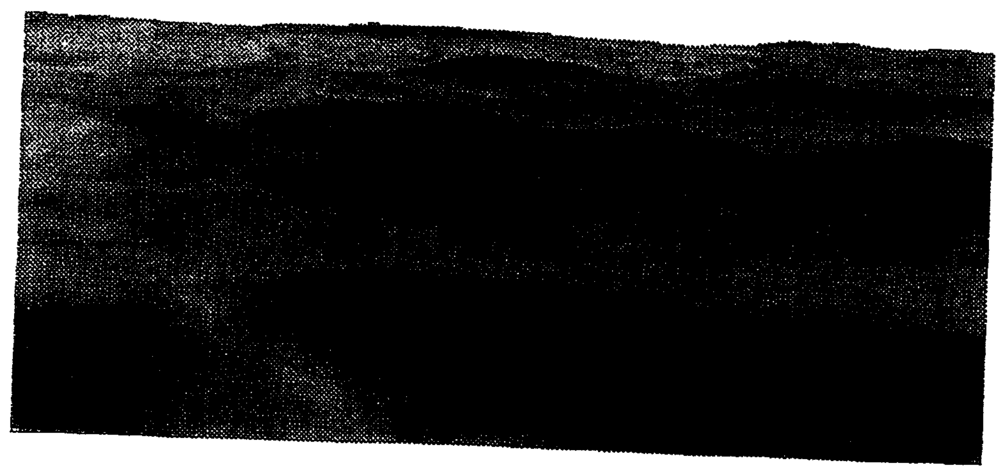

Elevation Data Rendered with Orthographic Image of Terrain

How Does TerraVision Work over a Network?

- ‡ Instead of copying the entire database into local storage, and then doing the visualization, TerraVision keeps only a small portion of the database in a local cache.
- ‡ It updates its local cache in real time by prefetching data depending on the user's viewpoint.

How does TerraVision use a Small Local Cache?

‡ The orthoimage mosaic is divided into small, equal-sized tiles.
- This way, the user can view a portion of the terrain using only a small number of "visible" tiles.
- Tiles are prefetched across the network (by predicting the user's path) and kept in a local cache until they are needed.

‡ The orthoimage mosaic is preprocessed into a multi-resolution hierarchy.
- Oblique views (encompassing a large area) can be rendered with a small number of tiles.
- Low-resolution tiles (covering large areas with relatively few pixels) can be permanently resident, so that a low-resolution image can always be rendered instantly.

TerraVision

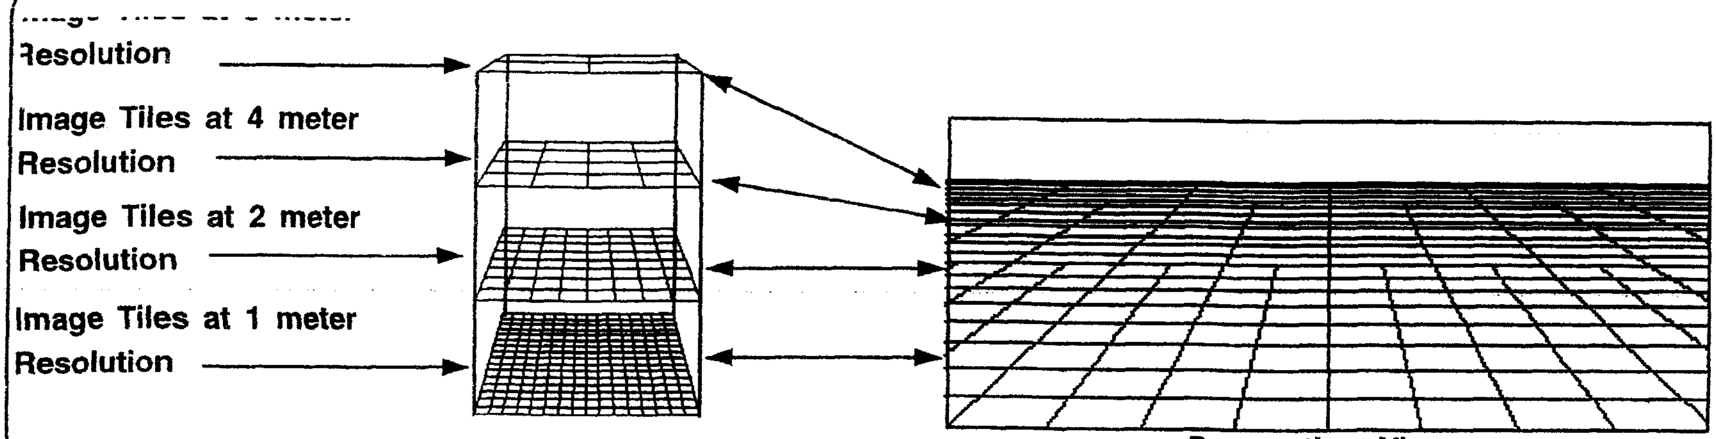

Different Resolution Tiles of Terrain Data

Overview of Placement and Retrieval of Tiles While Flying Over Terrain

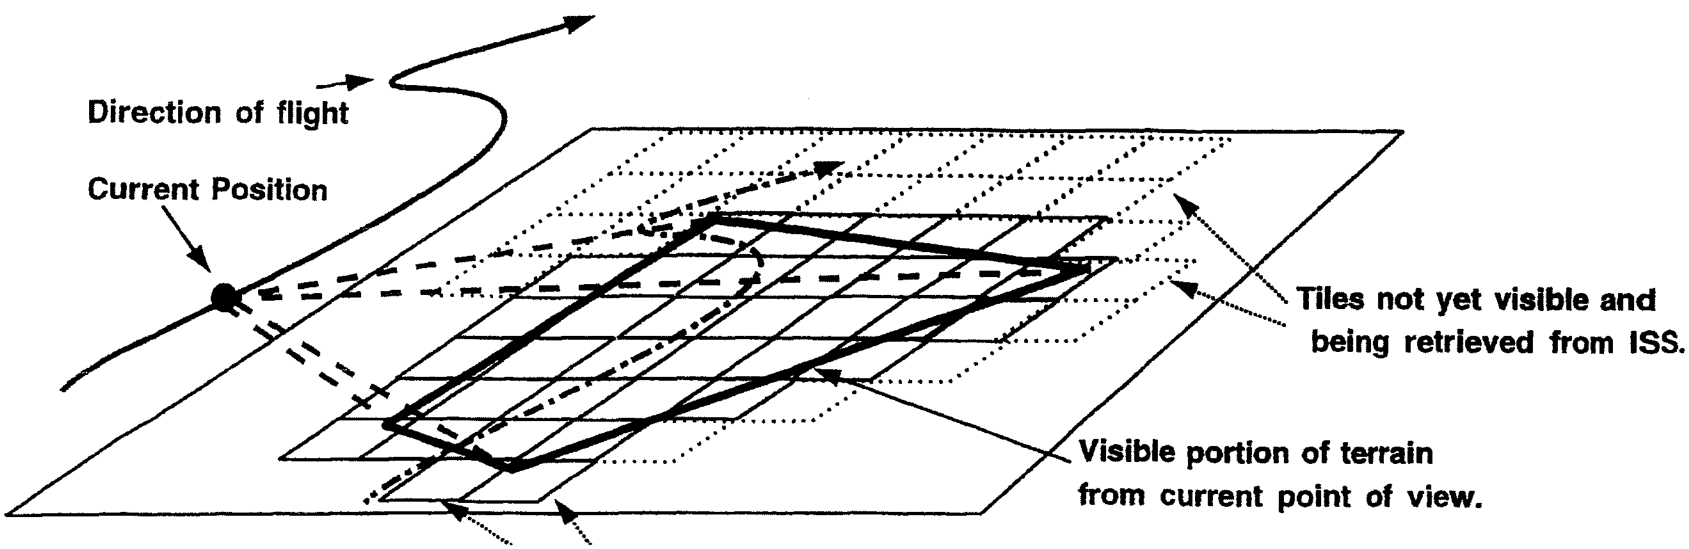

Contribution to the MAGIC Goals

‡ TerraVision is of direct interest to the Department of Defense because it can be used in situations requiring large data sets that are remotely created and maintained, such as

- Battlefield training
- 3D overview of new sites.

‡ TerraVision is a demanding application in which users can immediately see the impact of network congestion. It is therefore an excellent real-world application for testing the network.

Technical Significance

‡ The TerraVision implementation demonstrates that

- High-speed networks and data-distribution systems can indeed be used for real-time terrain visualization.
- Even very large terrain databases can be rendered in real time.

Implementation Challenges

- ‡ Creating a data representation that allows a large terrain database to be rendered with only a small amount of data
- ‡ Developing a search procedure that quickly finds the visible tiles in the hierarchy for any given viewpoint
- ‡ Developing a procedure that renders the terrain even when some of the data is not available in local memory
- ‡ Predicting a user's path so that data can be pre-fetched and a high image quality can be maintained.

Accomplishments to Date

‡ TerraVision (v1) has been implemented and demonstrated.
- Data rates of up to 100 mbits/second delivered to the application have been observed.
- Only one data set can be viewed per session.
- KU has implemented a remote viewing capability.

‡ TerraVision (v2) is ready for testing.
- Multiple data sets can be viewed per session.
- 3D controls and map views have been improved.
- Higher data rates and frame rates are anticipated.
- Remote control software (TeReVision) has been implemented, so that users on remote workstations can view and control TerraVision (based on KU work).

Lessons Learned

- ‡ Multithreaded, shared-memory processing appears to be the correct model for this type of application.
- ‡ Paging texture maps (tiles) into the graphics hardware is much slower than advertised, and is currently a bottleneck in the visualization process.

Script for TerraVision Videos

Yvan G. Leclerc
SRI

[TerraVision I title slide, 5 seconds]

TerraVision is a real time terrain visualization system developed at SRI International.

[HPCC slide, 13 seconds]

It is the test application of the ARPA-sponsored MAGIC project, one of several large-scale gigabit testbed projects. The underlying technology in TerraVision was derived in part from research carried out under the ARPA Image Understanding program.

[MAGIC participants, 17 seconds]

MAGIC, or the Multidimensional Applications Gigabit Internet Consortium, comprises a number of ARPA-sponsored and contributing participants.

The goal of the MAGIC project is to develop a gigabit-per-second ATM-based network and to address issues in high-speed networking, data storage, transport, and visualization.

[MAGIC map, 7 seconds]

The network currently connects Minneapolis; Sioux Falls, South Dakota; and Kansas City, Lawrence, and Fort Leavenworth in Kansas.

[fly through, 37 seconds]

TerraVision uses high-speed networks and distributed on-line storage systems to visualize multi-gigabyte-sized terrain databases, such as the one we are viewing, in a novel way.

Instead of copying the entire database to local storage before visualizing it, data is fetched across the network as it is needed while the user moves about the terrain. TerraVision is designed to accomodate the losses or delays that are inherent in the transmission of remotely stored data over a network.

The terrain data comprises an elevation model and a mosaic of pre-processed aerial images. It can be viewed in 3 dimensions, where the aerial images are draped over the terrain, as we have seen here.

[overhead view, panning and zooming, 11 seconds]

Or, it can be viewed in 2 dimension by panning and zooming over the terrain, rapidly moving from high-altitude panoramic views to low-altitude detailed views using push-button controls.

[cut to view of map, clicking several times, 5 seconds]

Another way to move about the terrain is to use a registered map on which the view area is outlined.

[cut to overhead view, while clicking on map, then perhaps to split window, with overhead and map]
[5 seconds]

The user can click anywhere on the map, instantly changing the point of view.

[turn on building models, 15 seconds]

The terrain data we are viewing are georeferenced. That is, the aerial images and elevation data have been precisely aligned with each other, and with a standard world coordinate system. What this means, for example, is that models of buildings can be accurately superimposed on the terrain.

[turn on GATS data, 22 seconds]

In addition, TerraVision can display a number of moving vehicles whose real-world positions have been obtained from sources such as portable GPS receivers. These vehicle positions can be directly transmitted to TerraVision for real-time viewing, or they can be recorded for future playback, which is what we see here.

Other types of mobile objects could also be viewed in real time using TerraVision.

[summary slide, 10 seconds]

In summary, TerraVision is a system for visualizing massive amounts of real georeferenced terrain data using high-speed networks and distributed on-line storage systems.

[advantages slide, 32 seconds]

Some of the advantages of this approach to visualization are:

o Visualization is not limited by the amount of workstation memory.

o Visualization is not limited by the size or access speed of local storage systems.

o Newly obtained data can be visualized simultaneously by multiple users at different locations, while communicating about common aspects of the data.

o A single geographically distributed database can be created and maintained by many experts, each responsible for one aspect of the composite information. Consistency of the data is therefore guaranteed.

[Acknowledgments, 30 seconds]

This research was sponsored in part by ARPA under contract F19627-92-C-0071.

Vehicle data (GATS) by Tom Gamet and David Rush, SRI International.

Original aerial images scanned and triangulated by the US Geological Survey, Menlo Park,CA.

Orthoimages processed using TerraForm on a CM-200 at MSCI by Yvan Leclerc, SRI International and Jay Feuquay, EROS Data Center.

The Image Server System was developed at Lawrence Berkeley Laboratory by Brian Tierney, Bill Johnston, Gary Hoo, Jason Lee, and Hanan Herzog.

[POC, 30 seconds]

For more information about TerraVision, contact:

Yvan Leclerc
Room EK252
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA, 94025

(415) 859-6153
leclerc@ai.sri.com

Or see URL:

http://www.ai.sri.com/aic/perception/projects/magic

[TerraVision II title slide, 7 seconds]

TerraVision is a real time terrain visualization system developed at SRI International as part of the ARPA-sponsored MAGIC project.

[2d view, moving from high to low altitude, 15 seconds]

To understand how TerraVision works, let's first look at the 2D overhead view. Note that from a high altitude, we can see a large area, but with little detail. Now as we move to a lower altitude, we see a smaller area, but with greater detail.

[turn on tile boundaries at high res, 21 seconds]

To see how this is done, let's see how the images are first processed and stored. The high resolution image mosaic, of which we are now seeing only a very small fraction, can contain as many as several billion pixels. It is divided into square tiles that are stored on a distributed on-line storage system called the Image Server System, or ISS, developed by Lawrence Berkeley Laboratory.

[ISS slide, 13 seconds]

The ISS comprises a number of workstations, each with a number of disks, all operating in parallel. As the user moves over the terrain, TerraVision requests tiles from the ISS and stores them in its local memory until they are needed for display.

[2d view, panning, 11 seconds]

To compensate for the delays inherent in requesting, transmitting, and receiving tiles, TerraVision anticipates the user's future viewpoint, and requests the tiles that cover that viewpoint ahead of time.

[2d view, zooming out, 20 seconds]

Now, as we zoom out, we see progressively more of the terrain. If we always used the high-resolution tiles for the display, TerraVision would eventually need all of the tiles in the remote database. Instead, the high-resolution image is preprocessed into a multi-resolution pyramid using a Connection Machine supercomputer located at the Minnesotta Supercomputer Center.

[2d view, changing resolution, but not altitude, with boundaries on]
[19 seconds]

At each level of the resolution pyramid, groups of 4 tiles from the next higher resolution are averaged down into a single tile. Consequently, each level of the

pyramid covers the entire terrain, but uses only 1/4 as many tiles as the previous level. The pyramid is built layer by layer until the entire terrain is represented by a single tile.

[panning and zooming, with tile boundaries on, 24 seconds]

By outlining the boundaries of the tiles, we can now see that as we pan and zoom over the terrain, TerraVision requests and uses only those tiles that it needs at the appropriate resolution.

But what if some of the tiles needed for a given view are not in local memory when they are needed for the display? This happens most often when we use the map interface, because TerraVision cannot predict where the user is moving, and hence cannot prefetch the required tiles.

[jump to new locations, with boundaries on, 33 seconds]

Notice that each time we click on the map, the image first seems out of focus, and then becomes clear. What's happening is that, when we first move to a new area, the high resolution tiles are not available in local memory, so TerraVision is forced to use lower-resolution tiles. At the same time as the display is being processed, TerraVision is requesting higher-resolution tiles from the server. As they arrive, TerraVision uses these higher-resolution tiles, and the image becomes progressively better focused. Of course, the faster the network, the quicker the image comes into focus.

[3d view, 11 seconds]

The same principles are used for 3D views. But here, we have the added complexity that, for oblique views, we need to use high resolution data in the foreground, and lower resolution data in the background.

[3d view of plane at oblique angle with tile boundaries] [or 3d view of plane looking straight on, then progressively tilting plane to show perspective effect]

[35 seconds]

To illustrate this point, let's look at a flat plane with only the tile boundaries being displayed. Note that as the plane recedes into the distance, the tiles become progressively smaller on the screen. If we were to let this process continue indefinitely, we would require a very large number of tiles to create this type of view. But, as you may have noticed, TerraVision uses progressively lower resolution tiles in the distance. This technique allows TerraVision to use approximately the same number of tiles for any view; and thus, tile requests can be made at an approximately steady rate as the user moves over the terrain, independently of the angle of view.

[clips from above, 25 seconds]

To summarize, TerraVision uses a resolution pyramid of tiles stored on the ISS. It anticipates the user's viewpoints and prefetches the tiles it requires. If the tiles are not all available in local memory, it uses the best ones it has, so that the user can always see something on the screen immediately. Finally, it uses only the appropriate resolution of tile, depending on the distance of the terrain from the user's viewpoint.

[Acknowledgments, 30 seconds]

This research was sponsored in part by ARPA under contract F19627-92-C-0071.

Vehicle data (GATS) by Tom Gamet and David Rush, SRI International.

Original aerial images scanned and triangulated by the US Geological Survey, Menlo Park,CA.

Orthoimages processed using TerraForm on a CM-200 at MSCI by Yvan Leclerc, SRI International and Jay Feuquay, EROS Data Center.

The Image Server System was developed at Lawrence Berkeley Laboratory by Brian Tierney, Bill Johnston, Gary Hoo, Jason Lee, and Hanan Herzog.

[POC, 30 seconds]

For more information about TerraVision, contact:

Yvan Leclerc
Room EK252
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA, 94025

(415) 859-6153
leclerc@ai.sri.com

Or see URL:

http://www.ai.sri.com/aic/perception/projects/magic